


ExCALIBUR

Finite Element Models: Complementary Activities I

M6.1

Abstract

This report describes work for ExCALIBUR project NEPTUNE at Milestone 6.1. It includes an application of the *Nektar++* finite-element software to the classic problem of vertical natural convection, with the aims of assessing the validity of the code in the context of equilibrium (non-kinetic) heat transport, and attempting to detect relevant challenges and issues. Numerical solutions for convective fields are obtained for a range of values of the Rayleigh number, the parameter governing the strength of the buoyancy force driving the convection; several regimes are observed, representing heat transport dominated by either diffusion or boundary-layer convective flow, and exhibiting complex time-dependent structures for larger Rayleigh number. The dependence of the rate of heat transfer on the Rayleigh number is examined. Also included in the report is an outline of the complementary experimental programme Smallab, and a brief mention of the work of the Oxford-led discussion group that focusses on the further development of fluid and kinetic models using finite elements. Appendices contain some technical details of the numerics, collect some work reproducing early studies of convection in water-filled tanks, and provide some initial estimates for the number of finite elements that may be mandated by the requirement to mesh accurately a complicated first wall geometry.



UKAEA REFERENCE AND APPROVAL SHEET

	Client Reference:		
	UKAEA Reference:	CD/EXCALIBUR-FMS/0051	
	Issue:	1.00	
	Date:	September 30, 2021	
Project Name: ExCALIBUR Fusion Modelling System			
	Name and Department	Signature	Date
Prepared By:	Ed Threlfall	N/A	September 30, 2021
	Wayne Arter	N/A	September 30, 2021
	BD		
Reviewed By:	Rob Akers		September 30, 2021
	Advanced Computing Dept. Manager		

1 Introduction

High-order methods, which offer the potential for exponential improvement in accuracy given polynomial increase in computation time, are seen as essential for maximizing the performance of next-generation software intended to run on forthcoming exascale machines. They are a natural fit to modern HPC due to the high arithmetic intensity of the algorithms - broadly, a relatively large number of floating point operations are applied per degree of freedom, thereby reducing the requirement for computer memory and, crucially, the expected demands on data throughput.

One specific example of such an approach is the spectral / *hp* finite element method [1], of which a modern C++ implementation is the software *Nektar++* [2] co-written and maintained by one of the NEPTUNE grantees. The *Nektar++* framework forms the basis for some of the core NEPTUNE proxyapps and is currently being augmented with capabilities to handle fusion-relevant problems ([3], [4]) and also extension to allow optimized operation on a wider range of computational architectures. This dependency motivates an examination of the software by UKAEA staff, included here in the application to vertical natural convection, a classic and historied problem which nevertheless has relevance to the issue of heat transport in the scrape-off layer (not to mention more direct applicability to fusion reactor components designed to handle large heat fluxes). In this preliminary work, results are obtained for two-dimensional convective flows. An attempt was made during the course of this work to detect any relevant issues with *Nektar* and one numerical problem is presented in an appendix. An account of planned experimental work concerning convective flows (Smallab) is included; the numerics are expected to complement this.

A brief account of the finite-element work of the Oxford-led discussion group (which meets regularly and includes the Oxford NEPTUNE grantee) is given, with reference to a growing codebase for gyrokinetic simulations written in the Julia language.

Any conclusions of the work with *Nektar++* are subject to the constraint placed on the framework by the need to mesh complicated first-wall geometries in a realistic simulation. An appendix contains brief estimates of the number of finite elements required in a realistic NEPTUNE simulation.

2 Vertical natural convection in *Nektar++*

One of the main problems of Earth-bound, controlled nuclear fusion is how to sustain a very large temperature gradient for a significant time, i.e. a problem of heat transport. It has long been known that classical (and neo-classical) transport calculations fail to account for the magnitude of the experimentally-observed rates of heat loss from magnetically-confined fusion plasmas, leading to the conclusion that turbulent and other phenomena are responsible for the bulk of the cross-field transport; the gradients in a tokamak (in temperature, pressure and magnetic field) are sufficiently extreme - e.g. approx. seven decades variation in temperature within the space of a few metres - that the system is a prime candidate for turbulent phenomena (which act universally to suppress large gradients). This invites study of simpler problems exhibiting qualitatively-similar behaviour.

One such example is vertical natural convection, in which an upright cavity has a hot side and a cold side, and is filled with a fluid that experiences a vertical buoyancy force when heated. The strength of this force then leads to different regimes of flow and heat transport. When the force is negligible, heat is transferred diffusively from the hot surface to the cold one, in rough analogy with classical transport. Greater buoyancy leads to fluid moving up the hot side, across the top of the cell, and down the cold side, creating a steady (laminar, or 'classical') circulation and transporting heat by advection. As the force is increased, the vertical flow becomes increasingly confined to a kinetic boundary layer near the vertical walls; likewise, the temperature gradients are increasingly restricted to a thermal boundary layer near the vertical walls. Further increasing the force, one finds that the sheared flow near the vertical boundaries leads to an instability of Kelvin-Helmholtz type that takes the form of waves that co-move with the convective flow, and which are chaotic in nature; these boundary effects also lead ultimately to complicated, also-chaotic, time-dependent flow structures in the bulk - turbulent flow.

The systems simulated in this study are purely two-dimensional, a fact which affects the results because three-dimensional flow instabilities are automatically excluded; also, there are important differences between turbulent flows in two dimensions and those in three. It might be noted that the two-dimensional property seems rather compatible with the fusion use case, with the turbulent and other transport phenomena - e.g. the well-known filaments - in a tokamak being strongly elongated along the magnetic field lines and so the dynamics of the scrape-off layer is, in this sense, two-dimensional.

In dimensionless form, the system of equations describing vertical natural convection is, for fluid velocity \mathbf{u} , temperature T and pressure p , for a system coordinatized such that the vertical, gravity-opposing direction is $\hat{\mathbf{y}}$,

$$\frac{1}{Pr} \left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} \right) = -\nabla p + Ra T \hat{\mathbf{y}} + \nabla^2 \mathbf{u}, \quad (1)$$

$$\left(\frac{\partial T}{\partial t} + \mathbf{u} \cdot \nabla T \right) = \nabla^2 T, \quad (2)$$

$$\nabla \cdot \mathbf{u} = 0. \quad (3)$$

The Boussinesq approximation is assumed, with the density fluctuation entering only as a temperature-dependent buoyancy force (hence the incompressibility condition) - for this reason, the system is

simulated using the incompressible Navier-Stokes solver of *Nektar++*. Dirichlet boundary conditions were used for the vertical wall temperatures ($T = 1$ on the left and $T = 0$ on the right) and zero Neumann conditions (i.e. no heat flux) on the horizontal walls (top and bottom). The velocity boundary conditions were no-slip / no-penetration and the pressure used a zero Neumann condition, for all walls.

The dimensionless parameters are the Rayleigh number, where β is the thermal expansion coefficient, ΔT the temperature difference across cavity width L ,

$$Ra = \frac{\beta g \Delta T L^3}{\kappa \nu} \quad (4)$$

and the Prandtl number: the ratio of kinematic viscosity ν to thermal diffusivity κ (a property of the fluid filling the cavity - small e.g. 0.01 for something like a liquid metal, rising to c. 10^5 for viscous oils),

$$Pr = \frac{\nu}{\kappa}. \quad (5)$$

The other relevant parameter is the aspect ratio $\lambda = \frac{L}{H}$ of the cavity (width L , height H).

The main output that is studied is the averaged Nusselt number, the physical meaning of which is that it measures the per-unit length heat flux through the cavity (i.e. going in to the left-hand side and out from the right-hand side) at any instant of time, defined as

$$Nu = \frac{\int_0^H -\nabla_x T(x_0, y) + \nabla_x T(x_1, y) dy}{\int_0^H -\nabla_x T(x_0, y) + \nabla_x T(x_1, y) dy|_{Ra=0}}. \quad (6)$$

Computational parameters enter as the mesh resolution (for simplicity, meshes comprising uniform squares are used for this study, with h denoting the size of a single element) and the global order p (a natural number) of the spectral elements, with the lowest allowed value ($p = 1$) meaning that linear intra-element basis functions are used, and greater p corresponding to higher-order Lagrange polynomials. For reference, a continuous Galerkin formulation was employed.

The open-source software *Nektar++* gives the ability to tackle research-grade problems in fluid mechanics on hardware ranging from a single PC to thousands-of-core HPC (note that, writing in the year 2021, it is quite possible to execute the supercomputer-class calculations of the 1980s and the early 1990s on a single modern PC - in the case of the results in this preliminary report, an Intel Core i7-10700KF 3.8GHz with 16 logical cores). Briefly, the software implements the spectral / hp element method for a range of problems. The convergence is, in ideal cases, expected to lead to L2 error scaling as $\sim h^{p+1}$ (see Appendix A for a demonstration). Note also that certain simulation data are efficiently extracted from *Nektar++* using *filters*; because filters for the quantities of interest needed in this investigation are not included in the source code, the software was modified in order to calculate the Nusselt number, and the temperature and location of the hottest point on the cold wall as time series. *Nektar++* contains a number of different solvers built on the same underlying finite-element framework; the one used here is the incompressible Navier-Stokes solver (details of the implementation can be found in [1]).

During the work presented in this report, a small anomaly was found in one part of the *Nektar++* incompressible Navier-Stokes algorithm (the pressure-solver option `VelocityCorrectionScheme`); an account is given in Appendix B.

2.1 Brief literature review

Historically, problems in convection have examined the output Nusselt number Nu as a function of the Rayleigh number Ra , typically concluding in a power-law relation of the type

$$Nu \sim Ra^\gamma \tag{7}$$

and proposing various values for the exponent γ , typically in the range $\frac{1}{4} - \frac{1}{2}$.

Early studies were undertaken by Elder in the 1960s. Numerical calculations were able at that time to access only low- Ra laminar flows ([5]); experiments with water-filled ($Pr = 7.0$) tanks allowed investigation of the phenomenology of more active flows, including a demonstration of the wall-wave instability ([6]). See Appendix C for some examples revisited using *Nektar++*.

More recent studies have focussed on air-filled cavities ($Pr = 0.71$) in the contexts of double-glazing and building cooling systems. The 1989 review [7] by Wright and Sullivan includes a survey of experimental results for various values of the aspect ratio in the range $\frac{1}{5}$ to $\frac{1}{110}$. The 2005 paper [8] by Xamán and collaborators obtained an approximate scaling exponent of 0.25 for a two-dimensional laminar convective model, rising to 0.3 when a three-dimensional turbulence model was included.

The current state-of-the-art appears to be the work of Lohse and collaborators [9]. An earlier work by Grossmann and Lohse [10], treating the related case of Rayleigh-Bénard convection, where the heating is applied from beneath, rather than from the side, included an analysis that revealed a range of physics-motivated scaling laws $Nu \sim Ra^\gamma Pr^\delta$ for various regimes differentiated by whether the boundary layer or the bulk dominates the global kinetic and thermal dissipation, and by which of the thermal or kinetic boundary layers is thicker. This theory led to an instructive ‘phase diagram’ of various flow regimes in the two-dimensional Pr, Ra parameter space, with their own values of γ and δ , the boundaries between regimes being determined by numerical experiment. It is notable that the same calculations are not possible in the case of vertical natural convection due to the lack of closure of the system of equations in that case. One might nevertheless expect something similar to exist in the case of vertical natural convection, viz. a collection of physically-motivated power laws that change as the parameter space in Ra and Pr is navigated; the authors of [9] state that they are working on this problem. A solution would rationalize the jumble of historic power-law fits.

It is to be emphasized that much of the literature, including the state-of-the-art mentioned above, is concerned with the three-dimensional case, whereas the results presented in this report are two-dimensional.

Also recently, the paper [11] has exposed an analogy between the Rayleigh-Bénard problem and the heat transport in a tokamak.

2.2 Air-filled cavity: simulation outputs

Matching the studies in [7], [8], and [9] a value of $Pr = 0.71$ was selected, this corresponding to the cavity being filled with air. Aspect ratios λ of $\frac{1}{2.5}$, $\frac{1}{5}$, and $\frac{1}{10}$ were studied. In the initial work, a two-dimensional mesh of uniform squares was used as the finite-element discretization, with 48 squares across the width of the cavity; thus $h = \frac{1}{48}$ and the cavity meshes comprised 5760, 11520, and 23040 elements. The lower- Ra simulations are easily simulated using as initial data the diffusive - $Ra = 0$ - solution (i.e. a uniform temperature gradient from the hot to the cold side, zero velocity, and no pressure gradient) though note it is advisable to add a small perturbation to this in order to avoid simulations getting 'stuck' on the diffusive solution. At higher Ra of c. 10^4 and above, this choice of initial condition leads to a violent transition to the new flow state, which is prone to numerical instability, and so the fields from a lower- Ra simulation (typically one decade less in Ra) are used as initial data. Note also that, for values of Ra below the critical value for the appearance of wall-wave instabilities, the problem is simply one of finding a stable steady state; for higher Ra , the system fluctuates about what is termed a 'quasi-steady state' in which statistical properties (e.g. the time-averaged value of Nu) are time-independent.

Plots of the temperature field in the steady state, or, in higher- Ra cases, the quasi-steady state, are shown in Figs.1, 2, and 3 for the three different aspect ratios. These demonstrate the transition from a diffusion-dominated solution (with temperature gradient in the horizontal direction) to a steady laminar (aka classical) convective cell solution (where the temperature gradient is increasingly in the vertical direction) and later to an unsteady flow which displays wall-wave instabilities, which are waves travelling up the hot wall and down the cold one (the properties of the boundary-layer waves are examined briefly in 2.3). It is obvious that the diffusive solution represents a strict minimum in heat transfer rate, and that increased Ra gives, in all cases studied here, increased heat transport. It is apparent that the smaller aspect ratio cavity retains its classical flow for greater values of Ra than the larger aspect ratio cases - this probably due to the fact that the near-wall velocities are greater for larger aspect ratio (Fig.5), since in that case the fluid is acted on by the buoyancy force over a greater distance; one also notes that the wall waves tend to appear first near the top of the hot wall and the bottom of the cold, where the vertical speed is expected to be maximal.

The data used in Figs.1-6 were taken from the most accurate simulation run for each choice of parameters (for the $\lambda = \frac{1}{2.5}$ cavity, $p = 3$ for $Ra = 10^5$, $p = 4$ for the higher- Ra ones, with the exception of the $Ra = 10^9$ case which was simulated using a $h = \frac{1}{6}$ mesh with $p = 15$, and checked by repeating the calculation for $p = 10 - 14$, after it was found that using high order in this way was more time-efficient); for the $\lambda = \frac{1}{5}$ cavity, $p = 3$ for $Ra \geq 10^5$; for the $\lambda = \frac{1}{10}$ cavity, $p = 2$ for $Ra \geq 10^5$; and the unquoted lower- Ra flows all used $p = 1$). See also 2.2.1.

The dependency of the averaged Nusselt number Nu on Ra is shown in Fig.4. This shows approximate power laws valid over much of the range; exponents of 0.2579, 0.2550, and 0.2701 were obtained over the Ra ranges $10^4 - 10^8$, $10^4 - 10^7$, and $10^4 - 10^7$ for aspect ratios $\frac{1}{2.5}$, $\frac{1}{5}$, and $\frac{1}{10}$ respectively. For quasi-steady cases, Nu was calculated using an average over the trend-free region of the Nusselt number time series. The reason why these numerics do not give exponents ≈ 0.3 , as found in [8], [10], is because the simulations presented here are two-dimensional, whereas those works refer to the three-dimensional case.

The dependency of the maximum mid-height vertical velocity on Ra is shown in Fig.5. The max-




Figure 1: Temperature fields for the $\lambda = \frac{1}{2.5}$ cavity (blue: cold; red: hot; range 0 to 1). Ra increases from 10^2 to 10^5 in decades from left to right across the upper row, and from 10^6 to 10^9 across the lower. Instability-driven wall waves appear in the $Ra = 10^8$ case near the top of the hot side and bottom of the cold side, occurring over a larger portion of the vertical walls in the $Ra = 10^9$ case.




Figure 2: Temperature fields for the $\lambda = \frac{1}{5}$ cavity. Ra increases from 10^2 to 10^7 in decades from left to right. Wall waves appear in the $Ra = 10^7$ case near the top of the hot side and bottom of the cold side.




Figure 3: Temperature fields for the $\lambda = \frac{1}{10}$ cavity. Ra increases from 10^2 to 10^7 in decades from left to right. Wall waves appear in the $Ra = 10^6$ case near the top of the hot side and bottom of the cold side; for this aspect ratio they are capable of developing into 'prominences' capable of advecting heat significant distances into the bulk (not seen for the smaller aspect ratio cases here).




Figure 4: Dependence of Nu on Ra for cavities with aspect ratios $\frac{1}{2.5}$, $\frac{1}{5}$, and $\frac{1}{10}$.

imum velocity appears rather well-fitted by a power law, with exponents 0.496, 0.490, and 0.515 being found over the Ra ranges $10^4 - 10^9$, $10^4 - 10^7$, and $10^4 - 10^6$ for aspect ratios $\frac{1}{2.5}$, $\frac{1}{5}$, and $\frac{1}{10}$ respectively. There is no significant deviation from these power laws at higher Ra . Note that the dependency is linear for smaller Ra due to the linear growth of the convective mode.

This discovery of the approximate relations $Nu \sim Ra^{\frac{1}{4}}$ and $v_{max} \sim Ra^{\frac{1}{2}}$ is consistent with the low- Pr Regime I_l for Rayleigh-Bénard convection found in [10], in their eqs (2.31) and (2.32) (note that v_{max} here stands as a proxy for the Reynolds number). This regime corresponds to viscous and thermal dissipations being dominated by the boundary layer, as opposed to the bulk, region, and the velocity boundary layer being less wide than the thermal boundary layer (i.e. small Pr) - though caution is advised applying these criteria to the present case as their derivation assumed the Rayleigh-Bénard case. It is also stated in that paper that the same relations have been found to be valid for low- Pr two-dimensional convection in [12].

The horizontal profiles of temperature and vertical velocity component (normalized to the maximum value) at the mid-height of the cavity are shown in Fig.6. It is evident that the temperature profile is linear in the diffusion-dominated regime, changing to something that is remarkably uniform except in the thermal boundary layers near the cavity walls, i.e. the gradients in the bulk are suppressed in the classical convective solutions; larger- Ra flows show significant disruption due to the non-steady phenomena. The vertical velocity profile varies from a sinusoid-like profile at low Ra to a flow dominated by boundary layer transport, before again becoming disrupted by non-steadiness. Large gradients in temperature and velocity become apparent near the boundary and these appear to increase monotonically with increasing Ra ; note that the near-boundary temperature exhibits apparently monotonic decrease with distance from the boundary, whereas the no-slip condition means that the vertical velocity rises to a sharp peak and then falls back. Note that time-averaging over the quasi-steady state can be performed to uncover the mean profiles in the unsteady cases, as done in [9]. The profiles in the figure are for the $\lambda = \frac{1}{2.5}$ cavity but




Figure 5: Dependence of maximum mid-height vertical flow velocity on Ra for cavities with aspect ratios $\frac{1}{2.5}$, $\frac{1}{5}$, and $\frac{1}{10}$.

are expected to be broadly similar for the other two cases, subject to rescalings in the Ra values demarcating the transitions between regimes.

2.2.1 Air-filled cavity: convergence, and pathways to convergence, and stability

The calculations ultimately become limited in terms of accessible Ra by the numerical resolution needed to simulate the velocity and temperature gradients: high- Ra is thus a computational challenge. In order to achieve convergence using the initial $h = \frac{1}{48}$ meshes, the element order was increased until the Nusselt number changed by less than a few percent. One point is that the convergence is harder to assess for the non-steady solutions since the averaged Nusselt number varies in time in the quasi-steady state (as an example, $\pm 10\%$ variation in Nu is seen for the $Ra = 10^9$ case). Note the the graphs plotted in order to obtain the power laws are rather forgiving in their response to errors (due to the logarithm). It was assumed that the convergence results for the smallest cavity ought to apply also to the larger cases (subject to rescalings of Ra e.g. it was assumed the flow corresponding to the onset of wall waves mandated a similar level of numerical precision in all cases).

The ability to change both h - and p -refinement seems currently to be under-explored in the convection literature (e.g. [9] uses a fully conservative fourth-order staggered finite difference scheme for the velocity field and the QUICK scheme to advect the temperature field, and the two-dimensional study in [8] uses a finite-volume method based on the method of Patankar). Initial results obtained during this work incline toward p -refinement: to cite the case of the $Ra = 10^7$ case for the

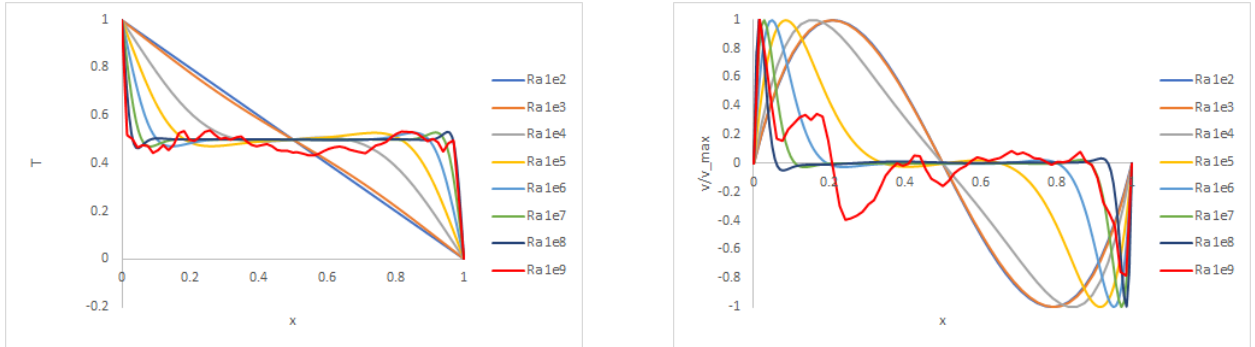


Figure 6: Mid-height horizontal temperature profile (left) and mid-height velocity normalized to maximum in profile (right) for $\lambda = \frac{1}{2.5}$ cavity.

p	$Nu (Ra = 10^5)$	$Nu (Ra = 10^6)$	$Nu (Ra = 10^7)$	$Nu (Ra = 10^8)$
1	4.123	7.279	12.286	18.231
2	4.173	7.601	14.222	26.613
3	4.167	7.523	13.711	25.222
4	—	7.554	—	—

Table 1: Convergence of Nu values with increasing element order p , for various Ra and $\lambda = \frac{1}{2.5}$ cavity / $h = \frac{1}{48}$ mesh.

$\lambda = \frac{1}{2.5}$ cavity, the $h = \frac{1}{48}$ $p = 3$ case (note $(p + 1)^2$ degrees of freedom (DOFs) per element) using $48 \times 120 \times 4^2 = 92,160$ DOFs and $\Delta t = 2 \times 10^{-6}$ took 3392s to propagate for 0.1 time units, while the apparently as-accurate $h = \frac{1}{6}$ $p = 11$ case (see Tables 1, 2) needed 796s using $6 \times 15 \times 12^2 = 12,960$ DOFs with $\Delta t = 5 \times 10^{-6}$ to evolve for the same amount of time (executed on the same hardware). The advantage here stems from the very significant reduction in the number of degrees of freedom needed to achieve a given level of accuracy, which is really the beauty of high-order methods. Note however that the simple structure here has no constraints on the level of h -refinement arising from the requirement to mesh a complicated surface accurately (see D.1 for some initial estimates in this direction). Note also that the question of performance-optimal discretization depends on machine architecture: in a massively-parallel system, it is likely that a decent number of elements will need to be associated with each compute core for efficient operation, and one certainly cannot have fewer than one element per core.

It may be of benefit to refine selectively the mesh near the wall, as, at least for the laminar flows, the largest velocity and temperature gradients are to be found there. For larger Ra , the existence of complicated time-dependent structures in the bulk (see the right-most case in Fig.1) probably means that a fine grid is needed everywhere. Note that it may be possible to increase element order p locally as well as performing local h -refinement.

It is also the case that the simulations encounter numerical instability that grows worse as Ra is increased (though this can be mitigated by matching the initial data closely to the quasi-steady state, where possible - as mentioned above). The precise determination of the maximum stable time step size is complicated and problem-dependent (see, for example, §6.3.1 of [1]) for these types of simulation (for purely linear systems the question is much easier, and indeed analytic

p	$Nu (Ra = 10^7)$	run time / s	Δt
3	13.936	46	5×10^{-6}
4	16.405	52	5×10^{-6}
5	15.286	88	5×10^{-6}
6	14.024	94	5×10^{-6}
7	13.509	109	5×10^{-6}
8	13.461	360	1.25×10^{-6}
9	13.557	467	1.25×10^{-6}
10	13.657	678	1.25×10^{-6}
11	13.716	796	1.25×10^{-6}
12	13.738	1198	5×10^{-7}
13	13.750	1418	5×10^{-7}

Table 2: Convergence of Nu values with increasing element order p , for $Ra = 10^7$ and $\lambda = \frac{1}{2.5}$ cavity / $h = \frac{1}{6}$ mesh. Also included are wall-clock run times in seconds (these are not definitive, as the time step size Δt was likely smaller than necessary, especially for the lower orders (< 8); also, some set-up time was included in the timings).

Ra	$\Delta t (p = 1)$	$\Delta t (p = 2)$	$\Delta t (p = 3)$	$\Delta t (p = 4)$
10^5	10^{-5}	10^{-5}	10^{-5}	—
10^6	10^{-5}	10^{-5}	5×10^{-6}	5×10^{-6}
10^7	10^{-5}	2×10^{-6}	2×10^{-6}	1.5×10^{-6}
10^8	2×10^{-6}	10^{-6}	5×10^{-7}	2.5×10^{-7}

Table 3: Table of stable time step sizes used for various Ra and $\lambda = \frac{1}{2.5}$ cavity / $h = \frac{1}{48}$ mesh.

for uniform computational meshes), and *Nektar++* seems to have no facility for obtaining even an estimate of a stable time step size - it is, to the author's knowledge, currently a case of trial and error. It might be noted that, from the aforementioned reference, the maximum stable time step size is expected to scale linearly with the size of the finite elements and inversely with the maximum local velocity (note the approximate scaling law $v_{max} \sim Ra^{\frac{1}{2}}$ found in the preceding subsection), with an additional penalty factor depending on the element order p . Some empirically-found stable time step sizes used in the simulations herein are presented in Table 3 (these represent lower bounds as a detailed search for the boundary between stable and unstable was not performed; also, 'stable' here means only that simulations run with these time step sizes showed no evidence of instability).

The question of performance is further muddled by the number of iterations in part of the solver being reduced as the time step size is reduced - thus a simulation with a larger number of time steps may potentially execute more quickly than one with fewer.

2.3 Transient time series and extreme event characterization

The time-dependent convective instabilities have implications for the time-variation in the global heat transport and also the variation in the local amount of heat transfer as a function of position

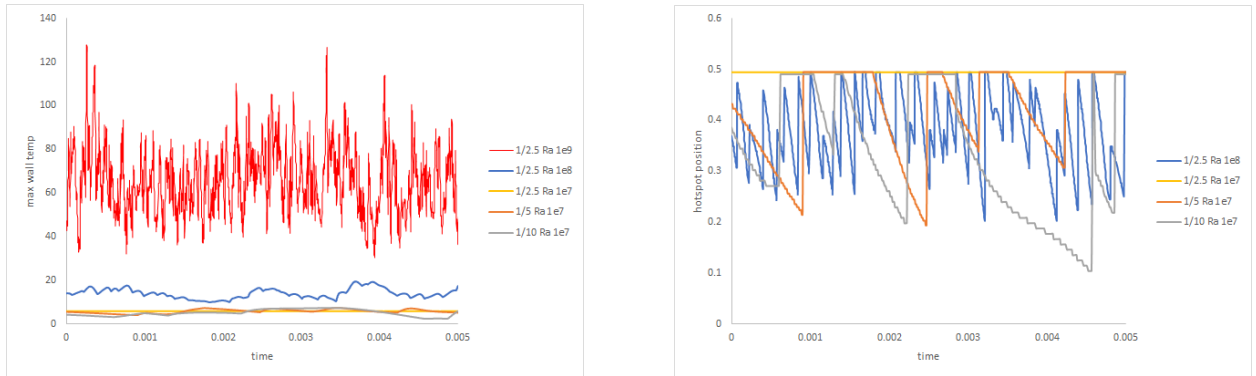


Figure 7: Transient behaviour associated to boundary-layer waves. The left-hand-side plot shows the maximum temperature at a small displacement from the cold-side wall and the right-hand-side one shows the vertical position of the hottest point near the cold wall, both as time series.

on the walls. This is of interest in the sense of study of ‘extreme events’ - the amount of local heating determines whether damage might result to a heat-absorbing engineering structure such as a tokamak first wall. The two-dimensional waves also have something of the character of the ‘filaments’ that have in recent years been observed near the outer boundary of tokamak plasmas.

To this purpose, a custom filter was added to *Nektar++*, designed to output the maximum near-wall temperature over a specified portion of the cavity side at regular time-intervals, and to output similarly the location of such a ‘hot spot’. Technically, the filter finds the maximum value of the normal derivative of the temperature field at the boundary (because of the fixed T values on the vertical boundaries, the temperature very close to the boundary is approximated as the boundary value minus $\epsilon \nabla_n T$ where ∇_n represents the outward normal derivative and ϵ a small spatial displacement), and note that this is in spirit a maximum value of the ‘local’ version of the averaged Nusselt number used above. Of course, in the steady state, the resulting time series are featureless, but in the cases exhibiting the wall-wave instability the outputs track e.g. cooling blobs of fluid as they descend the cold side.

The outputs (Fig.7), which were recorded in all cases for the lower 0.5 units of the cold-side wall, show clearly an increasing degree of variability in the maximum near-wall temperature as Ra is increased. The gradients of the hotspot position graph on the right-hand side show also that the wall waves move with greater velocity (represented by the gradient of these time series) for larger Ra (the graphs show hotter parts of fluid descending, cooling as they go, and ultimately being supplanted by the next hot wave entering the top of the monitored region). Note that the discrete character of the position graph comes from the fact that the filter currently outputs only the position of the intra-element Lagrange basis node and not the precise maximum. Note also that the $Ra = 10^9$ time series was excluded from the right-hand plot as it is sufficiently chaotic to obscure the other data; suffice it to say that the motion is a great deal more rapid in this case.

It is also possible to obtain time series for the averaged Nusselt number.

One issue is that the filter described in this section is currently not compatible with MPI execution of *Nektar++* due to its use of code that sums across threads (this is clearly not a problem for additive quantities such as the Nusselt number integrals described in the preceding section).

These time series have been used as the basis for uncertainty quantification and as training data for a Gaussian process. Further work could characterize the effect on heat transfer of the ‘prominences’ seen in the larger aspect ratio cavity (Fig.3) as these are expected to give rise to significant local heating events.

2.4 Conclusion

The incompressible Navier-Stokes solver of *Nektar++* was used to obtain plausible results for the problem of two-dimensional vertical natural convection. The relation between the Rayleigh number control parameter Ra and the heat transport as quantified by the averaged Nusselt number Nu was obtained for regimes ranging from the diffusive up to and beyond the onset of wall-wave instability and two-dimensional turbulence, for cavities of aspect ratio $\frac{1}{2.5}$, $\frac{1}{5}$, and $\frac{1}{10}$. *Nektar++* filters capable of producing time series for Nu , as well as for diagnostics of the near-wall dynamics, were demonstrated.

The main useful conclusion of this work is the indication that p -refinement is likely to lead to significant efficiency gains if used in preference to h -refinement. Another key issue is that of the maximum stable time step size which is one of the main criteria determining how much computational effort is entailed by a calculation and which was found to depend on the fineness of the discretization and on details of the flow through the Rayleigh number. Additionally, expertise in using *Nektar++* has been cultivated. This two-dimensional work prepares the path to more extensive simulations in more complicated geometries and in three dimensions, and is likely to prove a valuable guide to how best to organize future work on HPC.

3 Smallab experimental programme

In a cutting-edge technology programme such as fusion reactor development for STEP or DEMO, test facilities must be flexible to deal with issues which perhaps were not initially regarded as important and may require rapid resolution, for example in the light of new experimental data or during construction work. UKAEA’s Yorkshire Fusion Technology Facility FTF will be home to the CHIMERA test rig, to help develop and test materials and components that can cope with the conditions of high magnetic-field and heat-flux found inside a fusion power plant. FTF helps address critical STEP and DEMO design issues identified by Bachmann et al [13], viz. integrated design of breeding blanket and ancillary systems (their point ii) and power exhaust taking advantage of advanced divertor configurations (point iii). In support of CHIMERA, to verify scaling properties and provide physical insight, will be smaller, simpler table-top experiments such as Smallab.

New experiments, focussing on the key physics are justified because of better diagnostics enabled by greater data handling capability, advances in remote measurement techniques etc. For example, many of the key heat transport correlations are ripe for updating, as they date back 40 to 80 years, apply to a relatively restricted set of geometrical configurations, and come with little or no information as to their accuracy. Preliminary studies for NEPTUNE illustrate that there have been even greater advances in techniques for fitting models to data, in fields such as optimisation, uncertainty quantification and data assimilation, now often bracketed as Machine Learning, expected

to give more robust fitting with improved error estimates, relative to the old power-law scalings.

Increased computational power and the greater accuracy expected from NEPTUNE codes enables the use of more complex 2-D, 3-D and time-dependent physical models expected to be more robust under extrapolation to new situations and regimes. Nonetheless the presence of in particular turbulent behaviour means that there is still a need for cross-validation capabilities such as represented in particular by Smallab.

The concept is that Smallab experimental flows will be quick and easy to set up, not much more than the assembling of pipes, small vessels and heaters from stock. A typical example is the problem of convection in a slot with a heated sidewall, already investigated as part of NEPTUNE UQ activities [14] (and also in this report), a slightly more complex example is that of an annular pipe with a heated central core. Flow speeds of 5 m/s or less will be typical and similarly temperature differences limited to a few tens of degrees C.

The main outlay (near six-figure sum) has been to the LaVision company for a flow imaging system, specifically a MiniShaker + LED300 system. This itself may be rapidly set up and calibrated on timescales as little as 20 minutes. In the supplied diagnostic system, fluid flow is measured by Doppler effects in strong LED illumination. The raw data is processed by the DaVis 10 software supplied by LaVision to return flow vectors as high-resolution volumetric data sampled in time at up to 500 Hz. The DaVis 10 is capable of using knowledge of the equations of classical hydrodynamics to refine the flow-field information, which should be particularly interesting for comparison with *Nektar++* software. Any such comparison is promised to be made easier by the DaVis ability to output in a wide range of formats for visualisation and similar software on both Windows and linux. Should more complex analyses be required, the software is opensource C++ and therefore might be customised by skills widely available to project NEPTUNE.

It is to be noted that the numerics described in the preceding sections would have to be extended to three dimensions in order to apply to larger- Ra Smallab experiments, due to three-dimensional instabilities and turbulence.

4 Oxford-led discussion group meetings

The group have been meeting roughly fortnightly over Zoom. The meetings were chaired by Felix Parra until his departure for PPPL, and are now chaired by Michael Barnes. Other regular attendees are Michael Hardman and Javier Maurino-Alperovich (from the University of Oxford) and John Omotani, Sarah Newton, Stefan Mijin and Joseph Parker (from UKAEA), along with members of other NEPTUNE projects, Ben McMillan, Ben Dudson and David Dickinson.

The meetings were primary focused on Felix Parra and Michael Barnes presenting their respective progress with the development of the referent equations and their implementation, and to discuss challenges arising. The progress is presented in the project's technical reports. In addition, the project's source code is available at the GitHub repository [15] The code is presently 1+1D for space and velocity space parallel to the magnetic field line. In both dimensions, the user may choose a finite difference or Chebyshev (pseudo)spectral element discretization at run time. The code is written in Julia, making extensive use of external libraries (which in Julia are automatically managed and may be precompiled). The project has good software practices, making use of

GitHub issues and peer-reviewed pull requests. The code has unit and regression test suites that are automatically run on pull requests, as well as frameworks for code profiling and performance regression.

5 Summary

This report has provided an overview of preliminary application of *Nektar++* to a problem in thermal convection. It has identified that the dynamics of the system in question become more challenging for larger temperature differences (encoded in the dimensionless Rayleigh number) and the consequent greater rates of heat transfer (measured by the dimensionless averaged Nusselt number). The main challenges are the need to resolve steep gradients in the boundary regions and the fact that the maximum stable time step size decreases as the Rayleigh number is increased and also as the accuracy of the calculation is increased. Probably the main conclusion is the effectiveness of p -refinement in favour of h -refinement. In addition, attention was drawn to the restrictions on the maximum stable time step size imposed by the spatial discretization and also as a result of increasing the Rayleigh number (a more detailed study of this issue is clearly indicated). The complementary experimental programme *Smallab* was outlined with reference to its overlap with the *Nektar++* numerical simulations. In addition to this application of finite element methods, the Oxford-led discussion group continues to innovate in terms of applying finite element to more advanced referent models of the exotic fluids found in magnetically-confined fusion plasmas.

The numerical simulations outlined in this report used a very simple geometry; one eye is kept on the case of realistic first-wall geometries and the additional constraints applied by accurate meshing requirements; a small excursion in this direction is presented in Appendix D. Such geometrical considerations lead to the proliferation of elements and perhaps consequent moderation of the choice of element order, though more study needs to be done in order to uncover firm conclusions in this area.

The main author of this report would like to cite support received via the *Nektar++* mail list, from in particular Dr. Abhishek Kumar (Centre for Fluid and Complex Systems, Coventry University) who provided an example *Nektar++* session file for the Rayleigh-Bénard problem. The author would like to thank also the developers of *Nektar++* for useful discussions.

Acknowledgement

The support of the UK Meteorological Office and Strategic Priorities Fund is acknowledged.

A Convergence

In this appendix, the theoretical exponential convergence of the *Nektar++* software is demonstrated for the steady state of a classical convection problem solved on a fixed uniform 2D quadrilateral mesh.

For a smooth convecting solution in a cavity, the logarithmic error in the calculation (evaluated against a ‘converged’ high-order calculation for this simple smooth flow) was plotted against the element order. The error here is the global L2 error in the solution fields i.e. the squared error integrated over the entire domain, which is output automatically by *Nektar++*.

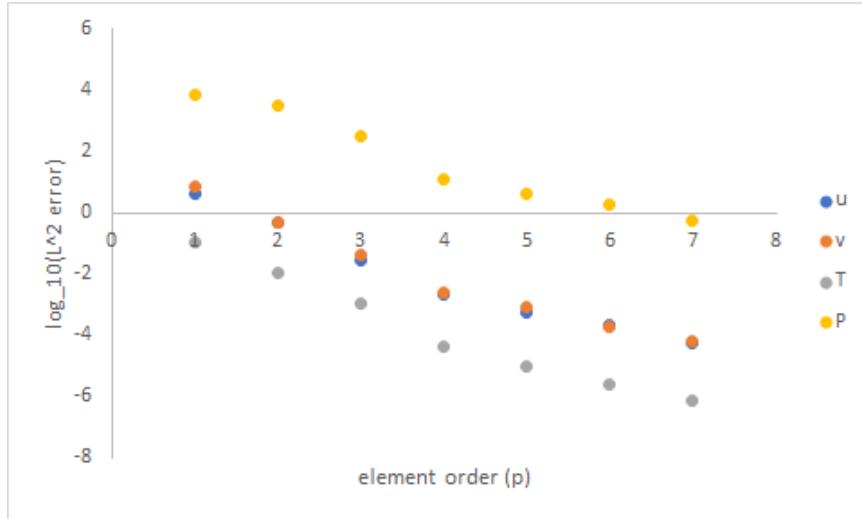


Figure 8: Error scaling with element order showing exponential convergence.

The results (Fig.8) show that this error decreases exponentially with increasing element order, in line with the theoretical promise for spectral / hp methods as applied to a smooth solution.

Note that this method cannot be applied directly to the higher- Ra , quasi-steady cases, as these are chaotic and only statistical properties of the flows are expected to remain consistent between simulations with different computational parameters (i.e. one cannot meaningfully compare individual instantaneous field values).

B *Nektar++* VelocityCorrectionScheme pressure anomaly

Simulations of smooth convective fields showed an obvious anomaly in the pressure field (left-hand-side plot in Fig.9): using the `VelocityCorrectionScheme` option for the incompressible Navier-Stokes solver resulted in a single boundary point having a pressure value fixed at zero, differing in a large amount from all neighbouring nodal values (note that the `VelocityCorrectionScheme` is a component of the incompressible Navier-Stokes algorithm that deals with decoupling the velocity and the pressure systems; the latter is solved using a Poisson equation with a homogeneous Neumann boundary condition).

In fact, the same problem as seen here can be created artificially using a Poisson equation where the Neumann boundary condition is inconsistent with the right-hand-side of the equation. The Neumann problem has a compatibility condition: if $\nabla^2 u = -f$ in a domain Ω and the boundary condition is $\alpha u + \beta \nabla_n u = g$ on $\partial\Omega$ then the divergence theorem gives the constraint $\int_{\partial\Omega} g + \int_{\Omega} f = 0$. As a test, this problem was set up in the two-dimensional advection-diffusion-reaction solver of *Nektar++* and it was seen that an inconsistent case gives a similar anomaly to that observed in the incompressible Navier-Stokes solver pressure field; also that the anomaly disappears if the right-hand-side term is offset to give a volume integral of zero. The hypothesis is that a similar inconsistency is present within the implementation of `VelocityCorrectionScheme`. It is




Figure 9: Left: pressure anomaly in the top-right domain corner using `VelocityCorrectionScheme` in `Nektar++`. Right: the same system solved using `VCSWeakPressure` instead, showing no issue.




Figure 10: Temperature contours (Fig.7(a) from [5]) on the left ($Ra = 2 \times 10^4$; $Pr = 1$; $\lambda = 1$ and note that the aspect ratio of the simulated cavity is actually unity, though not shown as such in the figure). The middle plot shows the `Nektar++` temperature profile at the end of the $p = 3$ simulation (steady state reached) and the rightmost image shows (if not very well) the 20×20 grid of finite elements.

presumably straightforward to perform the relevant integrals in the `Nektar++` framework to confirm this, though this is yet to be done. The need to solve this issue is tempered by the availability within `Nektar++` of an alternative to the `VelocityCorrectionScheme`, called `VCSWeakPressure`: this alternate does not show the anomaly.

C Elder's results

C.1 Elder's results

Early numerical work by Elder on classical convection, using $Pr = 1$, shows plots of the temperature field that concur well with the output of `Nektar++` (Fig.10). However, an accompanying graph of Nu as a function of Ra disagrees in the size of Nusselt number obtained for given Rayleigh number (Fig.11). The difference is probably due to a choice of boundary condition, given that the historic simulation used different degrees of freedom (namely velocity potential and vorticity).




FIGURE 9. Nusselt number N as a function of the Rayleigh number A .

Figure 11: Overlay of Elder's graph (from [5]) and Nektar++ results (blue spots) for the dependence of Nu on Ra .

Note that larger- Ra Nektar++ simulations using $Pr = 7.0$ (the value for water) show flow patterns (Fig.12) that compare quite well with water-tank experiments by the same author [6].

D Estimates for the number of finite elements in a NEPTUNE simulation

D.1 Estimate of number of elements required on surface meshing grounds

This subsection contains an estimate of the number of elements necessitated by surface meshing accuracy requirements found in the NEPTUNE grantee report [16].

The following assumptions are made:

- There are 500 structures to be meshed in the tokamak, each of which approximates a sphere or a box section as treated in [16].
- There are 20 elements needed to mesh away from the meshed surface to give a 3D mesh.

These assumptions together give an uplift of 10,000 above the number of elements needed to mesh the surface of a single structure such that there are no deviations between the mesh and the CAD greater than 0.1° .

Note that the box section structure treated in [16] may be a reasonable proxy for some of the components in a real tokamak e.g. the plasma-facing Faraday screen of the RF antennae used for ICRF heating and current drive, and also the ITER divertor support structure.

The number of elements needed to mesh an individual object is estimated from Tables 1 and 2 in [16] by interpolating or extrapolating the relation between the worst angle deviation and the number of needed elements to give the number of elements needed for a 0.1° worst deviation (the average angle deviation could also be considered). This interpolation / extrapolation was done




Figure 12: On the left is Fig.2(a) of [6] showing flow pattern for $Ra = 2.2 \times 10^7$ in a tank of water, visualized using dye injection. On the right is the quasi-steady state for the corresponding *Nektar++* simulation; this simulation used a uniform mesh 12 quadrilaterals across the horizontal axis and $p = 3$.

structure	order	num elts in mesh	num DOFs in computation
sphere	2	4.7×10^7	4.7×10^9
sphere	3	6.0×10^6	1.2×10^9
sphere	4	9.2×10^5	3.2×10^8
box	2	1.6×10^9	1.6×10^{11}
box	3	6.7×10^8	1.4×10^{11}
box	4	1.2×10^8	4.1×10^{10}
box	5	8.3×10^7	4.7×10^{10}

Table 4: Estimates of the number of elements required to achieve surface normal accuracy of 0.1° or better for the two structures tested in [16].

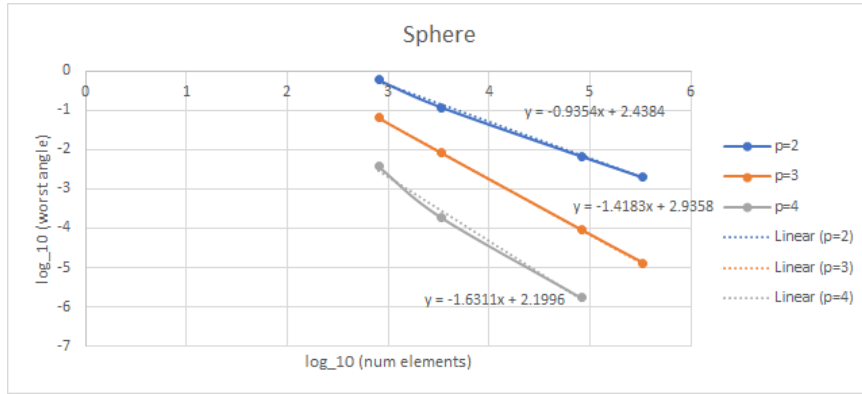


Figure 13: Log-linear fits for interpolating / extrapolating the number of elements needed to give $\theta_{max} = 0.1^\circ$, for sphere.

by constructing a plot of $\log_{10} \theta_{max}$ against $\log_{10} N_{elmt}$ where θ_{max} is the worst angle and N_{elmt} the number of elements in the surface mesh - the resulting plots are near linear (i.e. approximate power-law relations). Note that the 90° anomalies and also the highest-res $p = 4$ sphere mesh were excluded from these fits. For the sphere case, see Fig.13.

The results are shown in Table 4. Numbers include the uplift mentioned above. The DOF count assumes the 3D element computational order matches the geometrical order of the surface for all elements (and includes also an uplift of 10 to allow for multiple fields e.g. Maxwell, pressure, temperature ...) - for a single field the number of degrees of freedom is $N_{DOF} = \frac{1}{6}(p+3)(p+2)(p+1)$ for an element of order p i.e. assuming tetrahedral elements (other element geometries modify this formula but all scale $\sim p^3$).

Simple explanation - order-2 sphere - look at Fig.13 and see that the blue curve crosses $\log_{10} \theta_{max} = -1$ (i.e. 0.1°) at about 5,000 elements, then multiply by the uplift of 10,000 to get the number $4.7 * 10^7$; for the DOFs multiply this by 10 DOFs per elt at $p = 2$ and a factor of 10 to allow for all the different fields.

D.2 Order-of-magnitude estimates for tokamak edge modelling

This subsection contains some ballpark estimates of the computational feasibility of an anticipated full-scale NEPTUNE calculation.

Suppose plasma number density $n \approx 10^{18} \text{ m}^{-3}$. Order of magnitude dimensions are $L \approx 0.1 \text{ m}$ for SOL thickness, reactor minor and major radii say $a = 3 \text{ m}$ and $R_0 = 10 \text{ m}$, so volume of SOL $\approx 4\pi^2 a R_0 L \approx 100 \text{ m}^3$. Hence total number of electrons $\approx 10^{20}$.

Shortest timescale is inverse eB/m_e , the electron cyclotron frequency, where $e/m_e = 1.76 \times 10^{11} \text{ C kg}^{-1}$, and $B \approx 10 \text{ T}$, so $\tau_{Be} \approx 10^{-12} \text{ s}$. Hence number of particle-steps to evolve 1 s of physical time is $10^{20+12+1}$, and assuming 1000 flop per update, need 10^{36} flops. So to complete in 1 s on Exascale machine, only allowed to sample 1 particle-step in 10^{18} . (Unlikely to be adequate because of electrostatic and other effects.)

Suppose fluid instead, ie. representing electron distribution by first 3 moments. Electron temperature $T_e \approx 10 \text{ eV}$, thermal speed $V_{Te} \approx 10^6 \text{ m s}^{-1}$. Sample SOL at uniform 1 mm interval, number of sample-points $\approx 10^{11}$, timestep for explicit scheme $\approx 0.1 \times (10^{-3}/10^6)$ so number of sample-point updates is 10^{11+10} , assume 1000 flop each update, need 10^{24} flops.

Another way, suppose numerical problem is D -dimensional, need 1000 flop each sample update and allowed N_D samples per spatial dimension and N_D^2 in time. Then to update in 1 s, have $N_D^{D+2} \approx 10^{15}$. Thus if $D = 3$, $N_3 \approx 1000$, and $N_5 \approx 100$. *Accuracy controlled, unstructured, implicit fluid models should be possible.*

References

- [1] G. Karniadakis and S. Sherwin. *Spectral/hp element methods for computational fluid dynamics 2nd Ed.* Oxford University Press, 2005. <https://doi.org/10.1093/acprof:oso/9780198528692.001.0001>.
- [2] D. Moxey et al. Nektar++ website. <https://www.nektar.info>, 2020. Accessed: June 2020.
- [3] Nektar-diffusion proxyapp. <https://github.com/ExCALIBUR-NEPTUNE/nektar-diffusion>, 2021. Accessed: September 2021.
- [4] Nektar++ solver for Hasegawa-Wakatani equations. <https://github.com/ExCALIBUR-NEPTUNE/nektar-driftwave>, 2021. Accessed: September 2021.
- [5] J.W. Elder. Numerical experiments with free convection in a vertical slot. *Journal of Fluid Mechanics*, 24:823–843, 1965.
- [6] J.W. Elder. Turbulent free convection in a vertical slot. *Journal of Fluid Mechanics*, 23(1):99–111, 1965. <https://doi.org/10.1017/S0022112065001258>.
- [7] J.L. Wright and H.F. Sullivan. Natural convection in sealed glazing units: a review. *ASHRAE Transactions*, 95, 1989.

- [8] J. Xamán, G. Álvarez, L. Lira, and C. Estrada. Numerical study of heat transfer by laminar and turbulent natural convection in tall cavities of façade elements. *Energy and Buildings*, 37:787–794, 2005.
- [9] C.S. Ng, A. Ooi, D. Lohse, and D. Chung. Vertical natural convection: application of the unifying theory of thermal convection. *Journal of Fluid Mechanics*, 764:349–361, 2015.
- [10] S. Grossmann and D. Lohse. Scaling in thermal convection: a unifying theory. *Journal of Fluid Mechanics*, 407:27–56, 2000.
- [11] F. Wilczynski, D.W. Hughes, S. Van Loo, W. Arter, and F. Militello. Stability of scrape-off layer plasma: A modified Rayleigh–Bénard problem. *Physics of Plasmas*, 26(2):022510, 2019. <https://doi.org/10.1063/1.5064765>.
- [12] F.H. Busse and R.M. Clever. An asymptotic model of two-dimensional convection in the limit of low prandtl number. *Journal of Fluid Mechanics*, 102:75–83, 1981.
- [13] C. Bachmann, S. Ciattaglia, F. Cismondi, G. Federici, G. Keech, F. Maviglia, and M. Siccino. Critical Design Issues in DEMO and Solution Strategies. *Paper presented at SOFT 2018*, 2018.
- [14] E. Threlfall and W. Arter. Assessment of which UQ methods are required to make NEPTUNE software actionable. Technical Report CD/EXCALIBUR-FMS/0035-M2.4.1, UKAEA, 2021.
- [15] Moment_kinetics. https://github.com/mabarnes/moment_kinetics, 2021. Accessed: September 2021.
- [16] D. Moxey, C. Cantwell, and S. Sherwin. Surface mesh generation. Technical Report 2048465-TN-01-1, UKAEA Project Neptune, 2021.